

Our Ref.: 002717.P030

APPLICATION FOR UNITED STATES LETTERS PATENT

FOR

**Method and System to Aggregate
Multiple VLANs in a Metropolitan Area Network**

Inventor(s): **Michael Yip**

Prepared by:

BLAKELY SOKOLOFF TAYLOR & ZAFMAN, LLP
12400 Wilshire Boulevard, 7th Floor
Los Angeles, California 90025
(503) 684-6200

"Express Mail" label number EL034438983US

**Method and System to Aggregate
Multiple VLANs in a Metropolitan Area Network**

BACKGROUND OF THE INVENTION

5 1. Field of the Invention

The present invention relates to the field of metropolitan area network (MAN) topologies and internetwork communications technologies. In particular, the present invention relates aggregating multiple virtual local area networks (VLANs) into a MAN using VLAN identifier (VLAN ID) exchange.

10 2. Background Information and Description of Related Art

A metropolitan area network (MAN) is a network that interconnects users with computer resources in a geographic area or region larger than that covered by a local area network but smaller than the area covered by a wide area network. The term is typically applied to the interconnection of networks in a city into a single larger network (which may then also offer efficient connection to a wide area network). The amount of data traffic being sent over MANs is increasing at an exponential rate. This is due in part to the increasingly ubiquitous use of the Internet by consumers and businesses, as well as the increasingly bandwidth-intensive nature of the Internet applications that are being deployed.

An important aspect of MANs is the ability of MAN service providers to create virtual private network network connections (VPNs) across a single MAN infrastructure, referred to as a virtual metropolitan area network (VMAN). VMANs allow customers having multiple locations within a metropolitan area to transport private traffic, including virtual local area network (VLAN) traffic, over the shared single MAN.

However, the use of VMANs to handle traffic from multiple customers over a single MAN creates access and security issues. Therefore, it is important to segregate one customer from another so that there is no co-mingling of traffic.

In addition, customer traffic must be transported over the VMAN without
5 interfering with the customers' own higher-layer protocols such as DECnet, or private IP subnets. For example, the DECnet Phase IV protocol can cause problems when routed to a Layer 2 MAN because the DECnet protocol changes the media access control (MAC) address in the packet's Layer 2, or Data Link layer, header. Since duplicate MAC addresses are typically not allowed, VMAN service
10 providers end up managing DECnet streams by hand - something which neither the provider nor the customer wants.

The use of VMANs to handle traffic from multiple customers over a single MAN can also present scalability problems. For example, when switching VLANs based on the Institute for Electrical and Electronics Engineers (IEEE) 802.1Q
15 standard, the traditional system-wide upper limit of VLANs that can be handled while maintaining complete 802.1Q interoperability is 4,096. This may be insufficient for MAN service providers that need to provide network services to buildings having large numbers of individual customers connected to the network via traditional layer-2 VLANs.

20 Accordingly, a new approach is needed to securely manage traffic in a VMAN network architecture which does not interfere with higher level protocols and which is highly scalable.

SUMMARY

According to one aspect of the invention, a method and system is provided in which data packets from multiple customer VLANs are forwarded over a MAN using VLAN aggregation. A layer-2 switch located at the edge of the MAN connects the customer VLANs to the MAN. The edge switch aggregates multiple customer VLANs (the "sub-VLANs") into one provider VLAN (the "super-VLAN"). When a packet is forwarded from the sub-VLAN to the super-VLAN and vice versa, the edge switch uses modified bridge forwarding rules to exchange the customer-configured VLAN-IDs with the provider-configured VLAN-IDs before transporting the packet over the MAN via a VMAN layer-2 switch or MAN layer-3 router. The edge switch further uses modified bridge media access control (MAC) address learning rules to isolate one customer's traffic from another's (i.e. isolate one sub-VLAN's traffic from another sub-VLAN's traffic).

According to one aspect of the invention, apparatus are provided to carry out the above and other methods.

BRIEF DESCRIPTION OF DRAWINGS

The present invention will be described by way of exemplary embodiments, but not limitations, illustrated in the accompanying drawings in which like references
5 denote similar elements, and in which:

Figure 1 illustrates a diagram overview of a Metropolitan Area Network (MAN) configuration and aggregated VLANs in accordance with one embodiment of the present invention;

Figure 2a illustrates an example implementation of using VLAN aggregation
10 in a MAN in accordance with one embodiment of the present invention; and

Figure 2b illustrates the layer-2 edge switch of **Figure 2a** in further detail.

FIG. 1 is a diagram overview of a Metropolitan Area Network (MAN) configuration and aggregated VLANs in accordance with one embodiment of the present invention.

DETAILED DESCRIPTION OF THE INVENTION

In the following description various aspects of the present invention, a method and system in which data packets from multiple customer VLANs are forwarded over a MAN using VLAN aggregation, will be described. Specific details will be set forth in order to provide a thorough understanding of the present invention. However, it will be apparent to those skilled in the art that the present invention may be practiced with only some or all of the described aspects of the present invention, and with or without some or all of the specific details. In some instances, well known architectures, steps, and techniques have not been shown to avoid unnecessarily obscuring the present invention. For example, specific details are not provided as to whether the method and system is implemented in a router, server or gateway, as a software routine, hardware circuit, firmware, or a combination thereof.

Various operations will be described as multiple discrete steps performed in turn in a manner that is most helpful in understanding the present invention.

However, the order of description should not be construed as to imply that these operations are necessarily performed in the order they are presented, or even order dependent. Lastly, repeated usage of the phrase "in one embodiment" does not necessarily refer to the same embodiment, although it may.

A virtual local area network (VLAN) is a logical grouping of networked host computers on some other basis than the physical network location (e.g. customer, department, etc.). VLANs can be implemented in a number of different ways, depending on the network strategy. A prior art traditional layer-2 VLAN is based on a logical grouping of the layer-2 switch ports to which the hosts connect. Alternative prior art layer-2 VLANs define VLAN membership by the host's Media Access Control (MAC) layer address.

In some VLAN applications, data packets originating within a VLAN may carry a VLAN identification (VLAN ID) that can be used to provide intra-VLAN communication over a metropolitan area network (MAN) using existing layer-2 VLAN forwarding mechanisms. The existing VLAN forwarding mechanisms include both

5 proprietary and non-proprietary VLAN communication protocols. Currently, there is no single official standard protocol for communication of VLAN information. The method most commonly used is known as "frame-tagging." In frame-tagging, packets originating from a host belonging to a VLAN acquire a VLAN ID as they are switched onto a shared backbone network. The VLAN ID is what enables the receiving

10 switches to forward the packets intelligently by switching the packets to only those switches that are members of the same VLAN.

A non-proprietary VLAN communication protocol has been incorporated into the Institute for Electrical and Electronics Engineers (IEEE) 802.1Q standard, whereby the VLAN ID is part of the IEEE 802.1Q header inserted between the layer-2

15 Data Link header (i.e. the Media Access Control (MAC) header) and the frame's user data. This is also referred to as an 802.1Q tagged frame.

One way of using VLANs is to aggregate multiple VLANs hierarchically into a single "super-VLAN." The super-VLAN is then used to manage network traffic originating from the multiple VLANs, also referred to as sub-VLANs. An example of

20 aggregated VLAN architecture is described in a related application entitled "Method and System for VLAN Aggregation," which is assigned to Extreme Networks, Incorporated, the assignee of the present invention.

*Sub
a1*

In one aspect of the present invention, an aggregated VLAN architecture is used to create a virtual private network within a metropolitan area network (MAN),

25 also referred to as a virtual MAN (VMAN). Referring now to **Figure 1**, wherein a block diagram overview of a MAN configuration and aggregated VLANs in

accordance with one embodiment of the present invention is shown. As illustrated, a MAN 100 includes a VMAN layer-2 switch 102 and/or a layer-3 router 104.

Multiple customers 110, each having separate sub-VLANs, sub-VLAN 2 122, sub-VLAN 3 124, and sub-VLAN 4 126, are aggregated into super-VLAN 1 130 by a layer-2 edge switch 128 located at the edge of the MAN 100. As shown, the sub-VLANs 122, 124, and 126 are the customer-facing VLANs, whereas the super-VLAN 130 is the MAN-facing VLAN. The edge switch 128 forwards data packets originating from the customer sub-VLANs 122, 124, and 126 through the super-VLAN 1 130 and over the MAN 100 using the VMAN layer-2 switch 102 or the layer-3 router 104.

In one embodiment, data packets originating from the sub-VLANs may be tagged with a VLAN ID using an 802.1Q tag or other type of tagging scheme. In another embodiment, the data packets are not tagged. Either way, the edge switch 128 aggregates the sub-VLANs 122, 124, and 126 into the super-VLAN 130 by classifying the tagged or untagged packets according to the aggregated VLAN configuration. The aggregated VLAN configuration is typically pre-defined on the edge switch 128 by the MAN service provider, and is transparent to the individual customers whose VLANs are being aggregated. If the customer uses frame tagging, then the edge switch 128 simply verifies whether the VLAN ID specified in the data packet's 802.1Q tag is one of the configured VLAN IDs. If it is not one of the configured VLAN IDs according to the aggregated VLAN configuration, then the data packet is rejected. If the data packet is untagged, then the edge switch will assign a VLAN ID to the data packet, again according to the aggregated VLAN configuration.

Similarly, data packets originating from the super-VLAN 130 may be tagged with a VLAN ID using an 802.1Q tag or other type of tagging scheme. In another

embodiment, the data packets are not tagged. Either way, the edge switch **128** classifies the tagged or untagged packets according to the aggregated VLAN configuration. If the data packet is tagged, then the edge switch **128** simply verifies whether the VLAN ID specified in the data packet's 802.1Q tag is the configured super-VLAN's VLAN ID. If not, then the data packet is rejected. If the data packet is untagged, then the edge switch will assign a super-VLAN VLAN ID to the data packet according to the aggregated VLAN configuration.

In one embodiment, the aggregated VLAN configuration is composed of one sub-VLAN/VLAN ID for each customer and a single super-VLAN/VLAN ID.

However, other aggregated VLAN configurations may be employed without departing from the principles of the invention.

When the data packets are transported over the MAN **100** the MAN service provider must insure that the data packet is segregated from other customer traffic so as not to compromise the security of the packet or cause conflict with other customer traffic. For example, in residential buildings having layer-2 networks, it is important that each home not be able to communicate with each other, but only with the router of the MAN to which they are connected. In some cases, the segregation can be accomplished with a virtual MAN (VMAN) tunneling protocol, in which case the data packets are encapsulated with a VMAN ID tag before being forwarded over the MAN as described in related application "Method and System for VMAN Protocol Layer-2 Packet Re-encapsulation." However, for some MAN service providers, segregating traffic using a VMAN protocol may be insufficient to handle the volume of individual customers. One reason is that there are only a limited number of VMAN ID tags that can be assigned to data packets using a VMAN protocol; moreover, the VMAN switch at the core of the MAN can only switch up to 4096 individual VMANs. Another reason is that the VMAN protocol does not work with a

MAN core that is composed of layer-3 routers instead of layer-2 switches.

Therefore, a different approach to segregating traffic is employed in the method of the present invention. Two aspects of the method are the modified bridge MAC address learning rules and the modified bridge forwarding rules as described below.

5 Specifically, in one embodiment, once the data packets have been classified with the proper VLAN-ID, the edge switch **128** uses the modified bridge MAC address learning rules to isolate one customer's traffic from another customer's traffic. The modified bridge MAC address learning rules are derived as follows: when a new MAC address is learned from the sub-VLANs, the address is added in
10 both the sub-VLAN's and the super-VLAN's MAC-forwarding data base (FDB) table; when a new MAC address is learned from the super-VLAN, the address is added to all of the sub_VLANs' and the super-VLAN's MAC-FDB table. An example of FDB entries is shown in **Fig. 2b** in FDB Entries **275**, as detailed in an example implementation described below.

15 When a packet is forwarded from the sub-VLAN to the super-VLAN and vice versa, the edge switch further uses the modified bridge forwarding rules to exchange the customer-configured VLAN-IDs with the provider-configured VLAN-IDs before transporting the packet over the MAN **100** via the VMAN layer-2 switch **102** or MAN layer-3 router **104**. In this way, the customer's layer-2 configuration is
20 transparent to the MAN, and the MAN provider can achieve greater scalability.

 Specifically, in one embodiment, when a data packet is forwarded from the sub-VLAN to the super-VLAN, the original VLAN-ID of the sub-VLAN is exchanged with the super-VLAN's VLAN-ID and forwarded out into the super-VLAN and over the MAN using the VMAN layer-2 switch or layer-3 router. Conversely, when a
25 packet is forwarded from the super-VLAN to the sub-VLAN, the original VLAN-ID of the super-VLAN is exchanged with the sub-VLAN's VLAN ID and forwarded out into

the sub-VLAN and to the customer. The super-VLAN's original VLAN ID is configured by the MAN service provider, and the sub-VLANs' original VLAN IDs are configured by the customers.

When using the modified bridge forwarding rules to exchange the customer-
5 configured VLAN-IDs with the provider-configured VLAN-IDs, the edge switch 128 may obtain the "original VLAN ID" either directly from the data packet, as in the case of tagged data packets, or from the classification result (i.e. the VLAN ID assignment) internal to the switch, as in the case of untagged data packets.

As can be seen from the foregoing description, the illustrated embodiment of
10 the present invention makes it possible to create separate customer and provider domains for data packets transported over a single MAN 100. The customer domain is preserved by the edge switch 128 by using the modified bridge forwarding rules to exchange the customer-configured VLAN-IDs with the provider-configured VLAN-IDs before transporting the packet over the MAN 100 and vice versa. The
15 ability to create different VLAN aggregation configurations gives the MAN service provider a opportunity to create VMANs for their own use that can transport traffic from a greater number of customers than are transported with other methods, and without interfering with the flow of customer traffic.

For example, provider VMANs created using the VLAN aggregation methods
20 of the present invention can be used to connect customers to third-party services such as Internet Service Providers (ISPs) or Application Service Providers (ASPs). Since the connections to the ISPs and ASPs are transported over a VMAN, the customers can easily switch ISPs without disrupting their service. MAN service providers can also use VMANs to consolidate traffic and centralize value-added
25 services like VPNs or managed firewalls. Rather than being forced to install and maintain equipment on or near the customer premises – an expensive, labor-

intensive task – providers can aggregate customer traffic for value-added services at a central office. This not only saves on the cost of providing administrative staff, but achieves better economies of scale and gives customers more reliable service.

Better economies of scale are achieved due in part to the fact that with VLAN

5 aggregation each edge switch can accommodate more than the 4096 customers dictated by the 802.1Q limit of 4096 VLANs in a single VMAN. This is especially important since each VMAN core switch can only set up and manage up to 4096 VMANs. Thus, the more customers that can be accommodated in a single VMAN, the better the scalability.

10 In the illustrated embodiment, multiple customer VLANs are aggregated into a single super-VLAN. However, it should be understood that other types of networks may be aggregated without departing from the principles of or exceeding the scope of the present invention. Moreover, while the description of the embodiments of the present invention address the method and system as it applies
15 to use by a MAN service provider, it is appreciated by those of ordinary skill in the art that method is generally applicable to any network service provider that services multiple customers over any Internetworking application including, Local Area Networks (LANs), and Wide Area Networks (WANs).

An example implementation of one embodiment of the method and system of
20 the present invention will now be described. With reference to **Figure 2a**, customers in Building A **210** are connected to a MAN layer-3 router network **200** via Router 1 **215**, customers in Building B **220** are connected via Router 2 **225**, and customers in Building C **235** are connected via Router 3 **235**. As shown, the Building C network **240** is comprised of a layer-2 switch **245** (i.e. the edge switch)
25 that connects sub-VLANs VLAN 2 **255**, VLAN 3 **260**, and VLAN 4 **265** to super-VLAN VLAN1 **250**. The layer-3 Router 3 **235** receives from the layer-2 switch **245**

an 802.1Q tagged frame specifying a data packet with a .1Q tag VLAN ID = VLAN1, belonging to the super-VLAN VLAN 1 250 in Building C 230. Router 3 235 forwards the data packet over the MAN 200 based on Internet Protocol (IP) routing rules. In the reverse direction, Router 3 235 receives from one of the other routers router 1 215, or router 2 220, a data packet destined for a customer in Building C 230. Router 3 235 forwards the data packet to super-VLAN 1 250 via the layer-2 switch 245. The layer-2 switch 245 exchanges the super-VLAN's VLAN ID (VLAN 1) with the customer-configured destination VLAN ID (e.g. VLAN 2), and then forwards the data packet to the destination customer in Building C 235.

10 An advantage of the VLAN aggregation is that it allows the MAN service provider to deploy multiple sites with layer 2 point-to-multipoint connections. The use of VLAN aggregation to create a VMAN centralizes routing, servers, and services while maintaining simple layer-2 distribution without compromising network security between customers. With reference again to the illustrated embodiment in 15 **Figure 2a**, the inter-customer isolation is apparent from the permissible inter-VLAN communications using VMAN aggregation 270 shown as follows: VLAN 2 cannot communicate with either VLAN 3 or 4, but can communicate with super-VLAN 1. Likewise, VLAN3 and VLAN4 can only communicate with super-VLAN 1, and not with each other or with VLAN 2.

20 With reference to **Figure 2b**, the permissible inter-VLAN communication of the illustrated embodiment is shown in further detail. The Layer-2 switch 245 is equipped with 4 ports, numbered 1-4. The modified bridge MAC address learning rules are stored in the forwarding data base (FDB) 275 of the layer-2 switch 245. The FDB 275 contains the MAC addresses that can be received on all of the layer-2 25 switch's ports, and uses the information to decide whether a data packet should be forwarded or filtered. Each FDB entry consists of the MAC address of the device,

an identifier for the port on which it was received, and an identifier for the VLAN to which the device belongs. As shown, devices A and B are connected to VLAN2 255 via port 2, device C is connected to VLAN 3 260 via port 3, and device D is connected to VLAN4 265 via port 4. Port 1 is used to connect to Router R 235 which, in turn, connects the layer-2 switch 245 to the MAN layer-3 router network 200.

Another advantage to using VLAN aggregation to create a VMAN is that it also allows each customer to independently use VLAN tagging schemes such as the standard IEEE 802.1Q frame tagging, thereby simplifying administrative issues because each customer requires no knowledge of the super-VLAN ID assigned to them. In addition, VLAN aggregation ensures that the service provider infrastructures retain highly scalable characteristics. For example, VLAN aggregation allows scaling beyond the traditional system-wide limit of 4,096 802.1Q VLANs while maintaining complete 802.1Q interoperability. By exchanging the customer-configured VLAN 802.1Q tags with provider-configured tags, these values are made only locally significant and thus tag re-use can occur in the network. Additionally, similar to the traffic isolation concepts of VLAN aggregation, individual customer VLANs may be aggregated and appear as one tagged super-VLAN within the core of the service provider's MAN while still providing inter-customer traffic isolation.

Accordingly, a novel method and system is described for using VLAN aggregation to forward data packets by a MAN switch connecting multiple customers across a single MAN infrastructure. From the foregoing description, those skilled in the art will recognize that many other variations of the present invention are possible. In particular, while the present invention has been described as being implemented in a network comprising one or more MAN switches, such as

edge switch **128**, a core VMAN layer-2 switch **102**, a core VMAN layer-3 router **104**, a super VLAN **130**, and customer VLANs **122**, **124**, and **126**, some of the logic may be distributed in other components of a network or internetwork application.

For example, embodiments of the invention may be represented as a

5 software product stored on a machine-accessible medium (also referred to as a computer-readable medium or a processor-readable medium). The machine-accessible medium may be any type of magnetic, optical, or electrical storage medium including a diskette, CD-ROM, memory device (volatile or non-volatile), or similar storage mechanism. The machine-accessible medium may contain various
10 sets of instructions, code sequences, configuration information, or other data. As an example, the procedures described herein for aggregating a customer VLAN into a super-VLAN by edge switch **128**, for performing the modified bridge MAC address learning rules, for performing the modified bridge forwarding rules, or forwarding an aggregated VLAN data packet by a MAN core switch **102** or router **104**, may be
15 stored on the machine-accessible medium. Those of ordinary skill in the art will appreciate that other instructions and operations necessary to implement the described invention may also be stored on the machine-accessible medium.

Thus, the present invention is not limited by the details described. Instead, the present invention can be practiced with modifications and alterations within the
20 spirit and scope of the appended claims.
